CDF Production Farm'06

P.Murat (CDF/FNAL CD), for the CDF offline group

- Scope of the talk: **recent improvements to the CDF Production Farm**
  - There have been minor changes to the hardware organization of the Farm
  - The major ones – to the infrastructure of the control software
- The changes:
  - driven by the accumulated operational experience
  - aimed in
    - reducing the load on the Farm operators
    - increasing the operational depth of CDF offline group
    - allowing CDF to do more using the Farm (**large scale ntuple production)**
- One of the drivers: migration from SAM v6 to SAM v7
- The arm transitioned from "state A" to "state B"
  - "B" is a much better state than "A"
  - **we want to present our experience and discuss how did we get there**

CDF Production Farm

Analysis Farm (CAF)	Production Farm (FNCDFSRV0)	Test Farm (FNCDF)

- Unified infrastructure of the CDF computing:
  - Central Analysis Facility (CAF) – a batch farm running Condor
  - "Large" Production Farm: yet another CAF (~320 CPUs, 320 Vms)
    - ~20% of the total onsite CPU capacity
  - Small test Farm("phase I"): very small CAF (~40 old nodes?)
  - "Borrow" CPU cycles from the analysis farm when needs peak up
- Achieved throughput:
  - ~ 25 Mln events/day (Oct'2005), about 4 times larger than the data taking rate
- Very high stability of the reconstruction code:
  - less than 1 crash per 10^8 events

3 steps of data processing

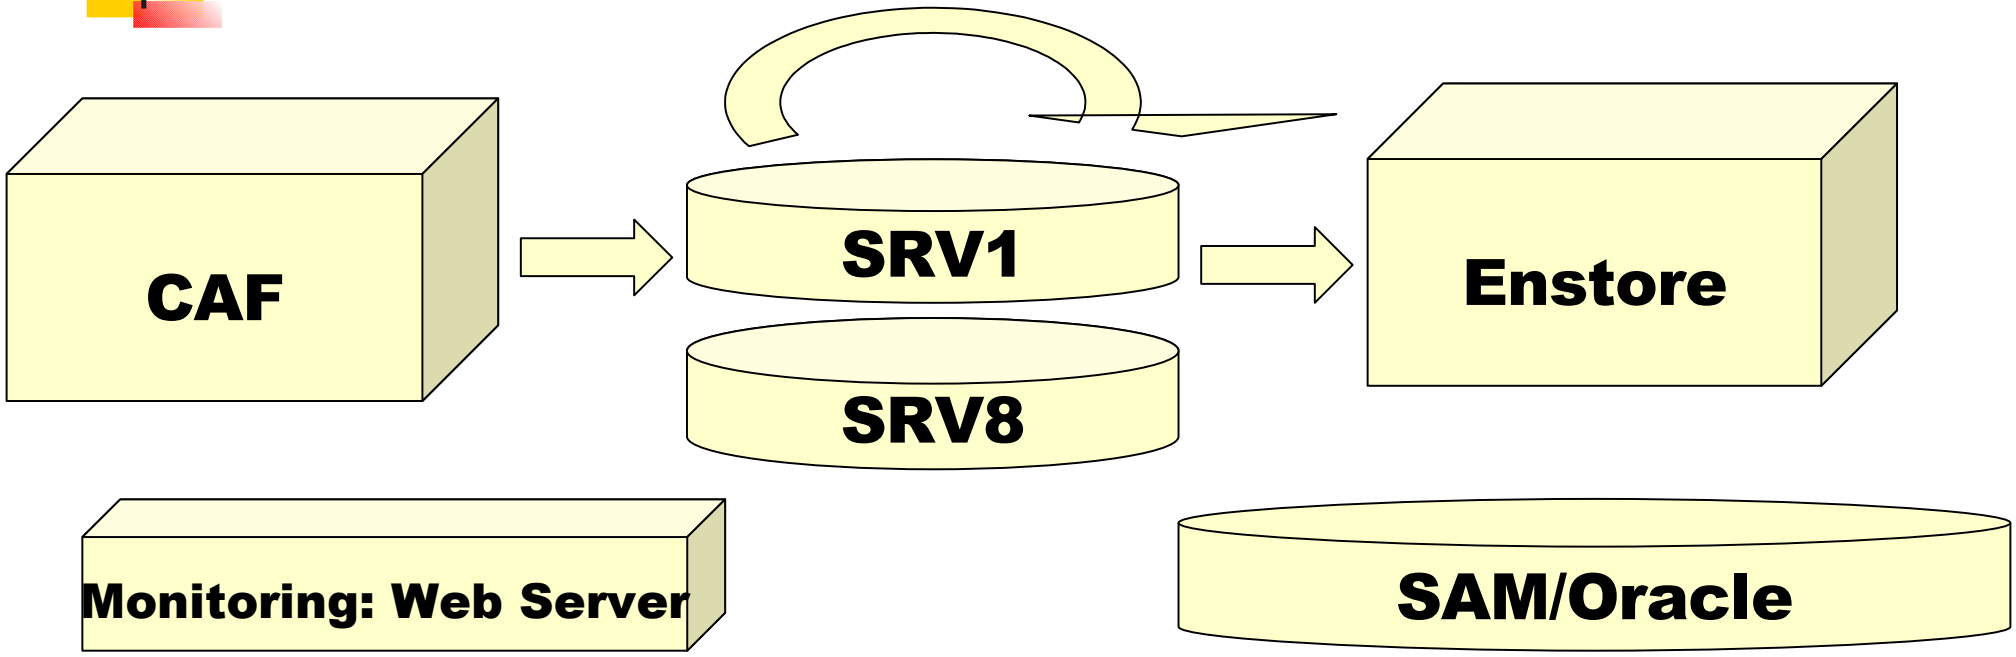

- 3 steps of the data processing:
  - Reconstruction: 7 input streams, ~50 output streams
  - Concatenation (8 servers, 2-7 Tbytes each)
  - Tape upload from the concatenation servers
- Bookkeeping: SAM
- Production Farm-specific monitoring: web-based, mostly static

The Farm: Operations

- Several issues related to the operations of the Farm:
  - **Why does it take 2 FTE's to operate the Farm ?**
  - **How can we reduce this number?**
- Primary task of the Farm – reconstruction of the data taken by CDF ("Production")
- Successfully implemented 1-pass data processing scheme:
  - The data reconstructed with the "final" calibrations become available within less then 2 months from the day when they were taken
- several related tasks handled by different CDF offline groups:
  - Beamline production
  - other offline calibrations (time-of-flight, calorimeter, timing readout in the CDF EM calorimeter
  - Offline data quality monitoring (calibration stream A)
- Needs very similar, non-overlapping parts relatively small, but sizable.
- Each group knows only their specific software, each having own personnel issues

Farm Transitions (I)

- Could all the tasks be done on the Production Farm using **the same** software infrastructure ? **What has to be changed for that?**
- Nov'2005: change of the ops personnel (Academia Cinica - > UNM)
  - The group which designed and implemented the Farm control software and kept responsibility for operating the Farm stepped down
- Specifics of transition:
  - Very little overlap between the old and the new Production Farm crews
  - Very strong team of offline ops managers
  - Next round of data processing started in the end of Dec'2005
- Several weeks of learning (in yellow – Farm CPU load around Christmas)

By the end of January we had all the answers implemented

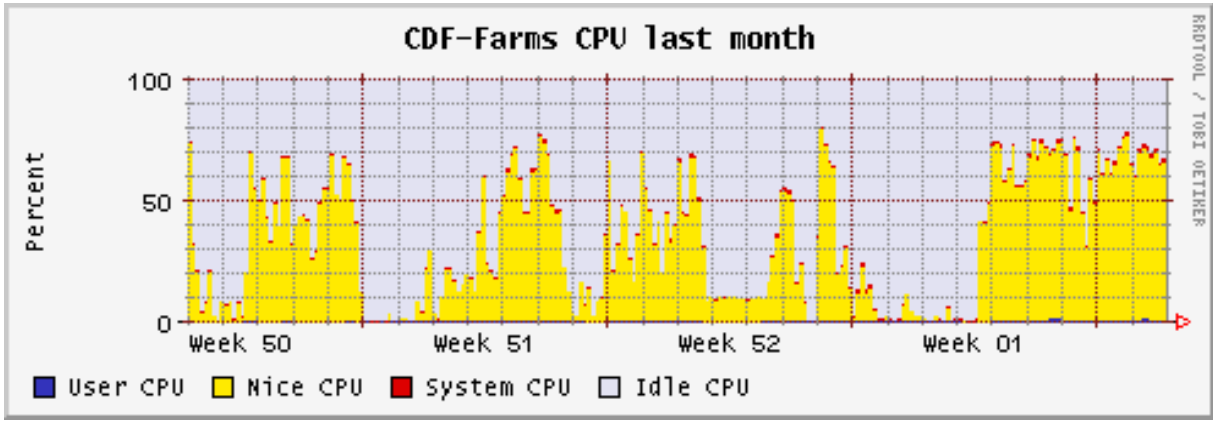

"Grand Unification" of the control software

- Control software cleaned up, put into CVS, use tagged versions
- Transitioned to SAM v7
- Logic of the control scripts generalized to allow
  - handling of all the calibration-related jobs
  - data quality monitoring
- Very important improvements in the automated error recovery procedures
- More robust logic of the job submission
- Moved to new scheme of Farm operations:
  - Common control software, different "back-ends" - executables doing the job
  - Well defined interfaces between the job control processes
- Next slides: few highlights

Better resource management

- Resource management has been done "by resource"
  - When a disk buffer was getting filled up, all the processes using it were suspended simultaneously and processing of the whole data stream stopped
- consequences:
  - more jobs submitted for other streams – the system goes off balance
  - Or less running jobs in total, efficiency of the CPU usage goes down
- **Implemented "by process" resource management**
  - When submission of Production job stops, concatenation and tape upload still continue such that the disk buffer gets cleaned up

Improved load balancing

- CDF Production Farm simultaneously processes up to 7 data streams
- Processed output of each stream goes to a separate concatenation fileserver
- Number of jobs submitted for different streams has to be balanced in order to balance the load on the concatenation fileservers
- Old logic:
  - number of running SAM projects for a given stream is less than N
  - Total number of running jobs was varying by a factor of up to ~4
- New logic**:**
  - **total number of used VM's for a given stream less than N, this keeps the load per stream constant within ~50%**

Improved efficiency of the CPU usage

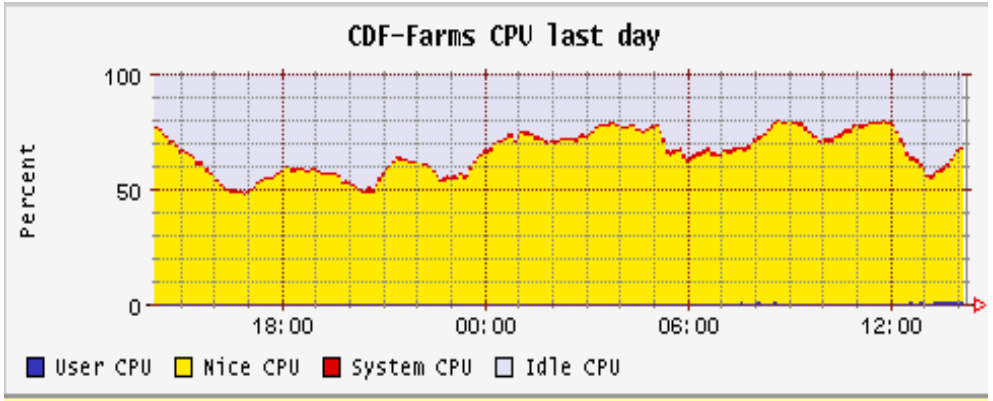

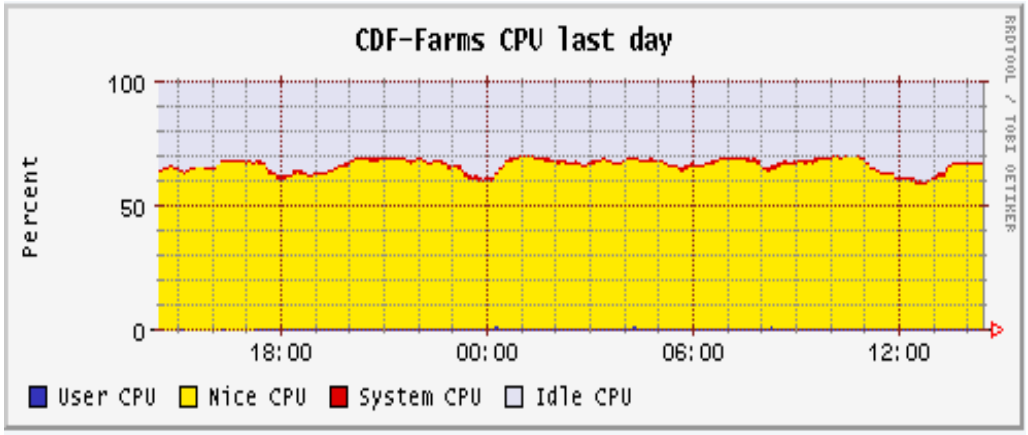

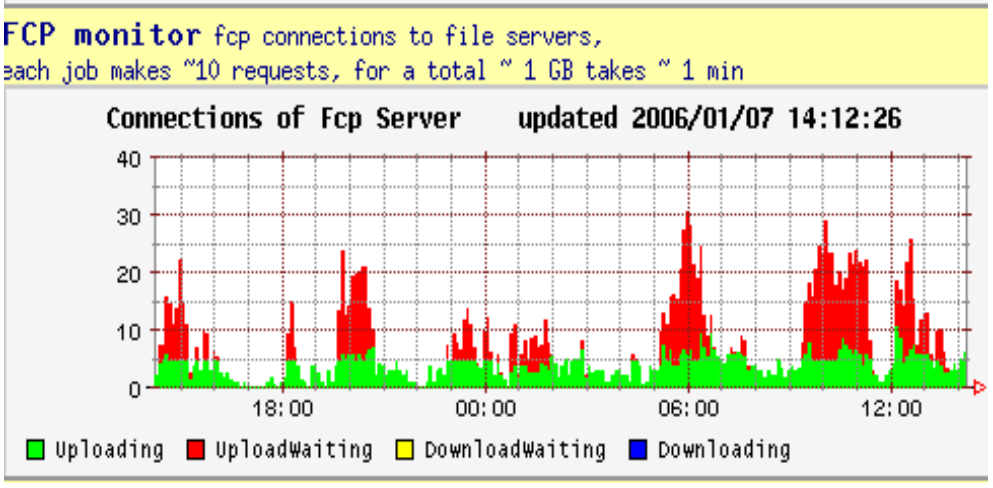

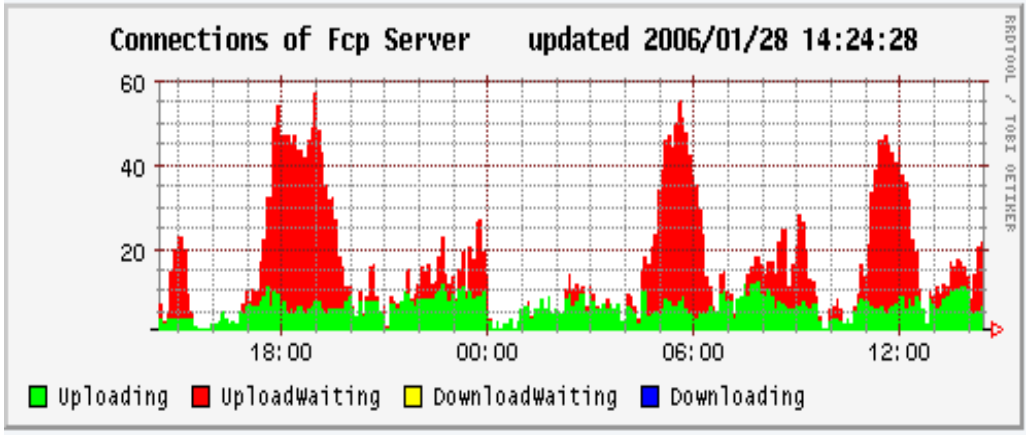

- Each node was running 2 "user" VM's, **increased the number up to 3**
- while the 1st job is waiting in FCP queue, the 2nd and the 3rd jobs are running

Jan 29: tests of the upgraded system

2 day-long test (~30M evens) on Jan 29:

- Load balancing worked
- no concatenation failures
- no tape writing failures
- no 'zombie' projects (although could've been)
- concatenation locks-ups expected, the source finally proven to be the expected one
- Preliminary conclusion: 100% success
- **Final conclusion: we did observe several failures, they , however, all of them have been automatically recovered**

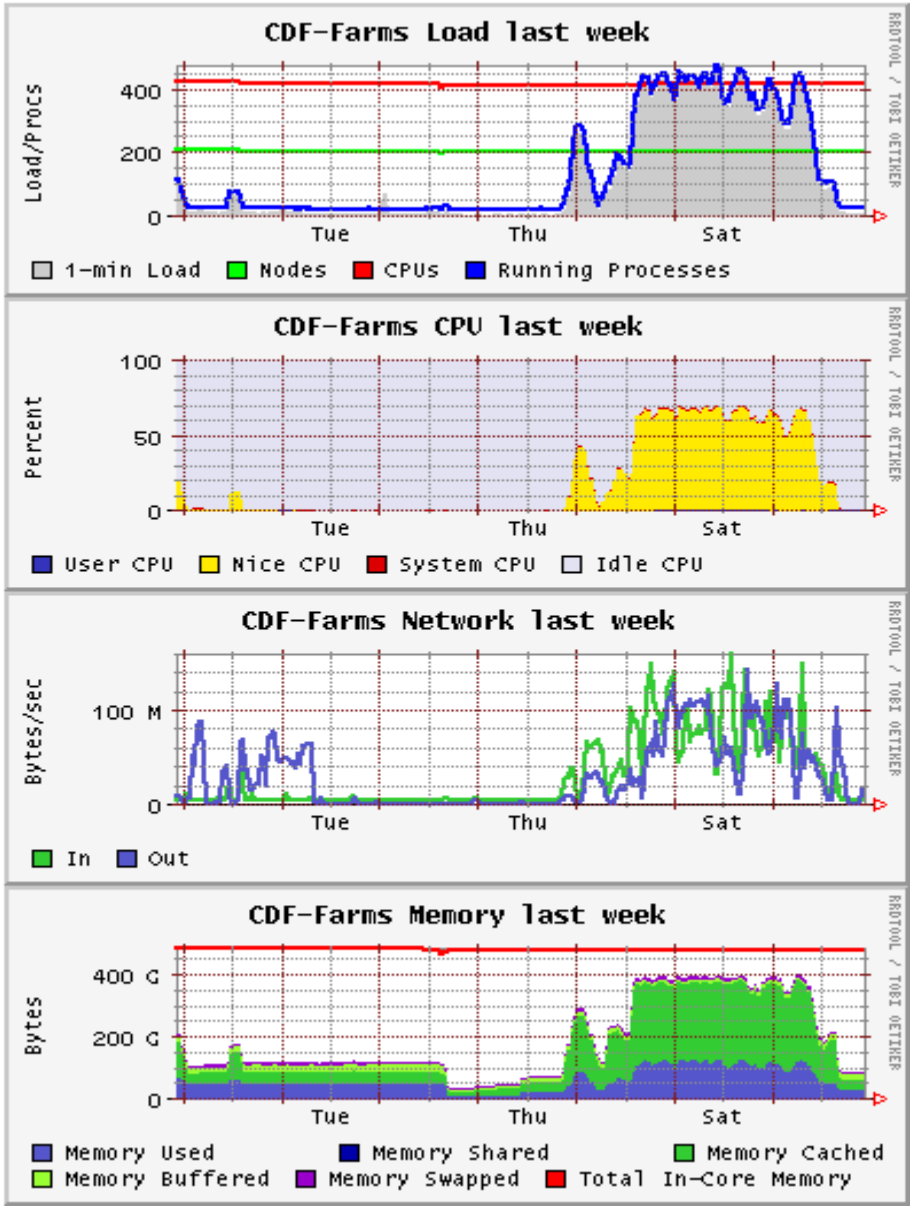

Large Scale Ntuple Production on the Farm

- Ntuple production needs were one of the major factors driving CDF estimates of the CPU requirements for FY'2006
- New considerations:
  - Ntupling code reached almost "production-level" maturity
  - Unified infrastructure of the Production Farm allows to handle ntuple production
- **This week CDF is holding workshop on ntuple production**
  - Power users, representing 3 major CDF ntuples will be trained to operate the Farm in ntuple production mode (**today – 5 trainees and 3 trainers !**)
  - Building operational depth of the Farms operations and getting the whole collaboration involved with it
- The plan: produce major ntuple datasets on the Farm
  - Expect very small extra delay: ntuples for the datasets written by Production are available in 2 weeks after the data are processed

Outlook

- Run IIB upgrade of the CDF Production Farm architecture complete
  - design of the hardware system and of the control software is solid and stable
- Test results indicate that we reached new level of stability:
  - expect the Farm operations to be taking less than 0.5 FTE
- This week starting new round of data processing (up to mid-January'06), expect to be done within 2 weeks
- Increasing operational depth of the Farms group - ~10 people including power users from the physics groups will be trained to operate the Farm
- Established the plan for producing the major ntuples using the Farm control infrastructure
  - will start executing it in the 2nd half of March
- **This will have major impact on the physics productivity of the CDF**